

CLAIMS

We claim:

1. A computer system for fault-tolerant distributed collaborative computing, the system
5 comprising:
 - a plurality of server computers connected to a plurality of client computers via a global-area computer network;
 - a high-speed direct connection link
10 connecting the plurality of server computers; and
 - a computer program executable by the server computers, wherein the computer program comprises computer instructions for:
 - conducting an on-line conference among
15 an arbitrary number of the client computers connected to an arbitrary number of the server computers via the global-area network and the high-speed direct connection link;
 - detecting a failure of one of the server
20 computers handling the on-line conference;
 - disconnecting the failed server computer from the on-line conference;
 - connecting another of the server computers to the conference; and
25 resuming the on-line conference.
2. The computer system of claim 1, wherein the computer program further comprises computer instructions for:
30
 - periodically replicating state information among processes executed by the server computers to conduct the on-line conference;

detecting a failure of one of the process;
spawning a new process on the server
computers; and
loading the replicated state information on
5 the new process.

3. The computer system of claim 2, wherein the
processes whose state is replicated maintain
information about the on-line conference.

10 4. The computer system of claim 2, wherein the
processes whose state is replicated handle
communications between one of the client computers and
one of the server computers.

15 5. The computer system of claim 2, wherein the
processes whose state is replicated control access to a
document shared among participants of the on-line
conference.

20 6. The computer system of claim 2, wherein the
processes whose state is replicated control execution
of an application shared among participants of the on-
line conference.

25 7. A method of operating a distributed
collaborative computing system comprising a plurality
of server computers, the method comprising:

conducting an on-line conference among an
30 arbitrary number of the client computers connected
to an arbitrary number of the server computers via

the global-area network and the high-speed direct connection link;

detecting a failure of one of the server computers handling the on-line conference;

5 disconnecting the failed server computer from the on-line conference;

 connecting another of the server computers to the conference; and

 resuming the on-line conference.

10

8. The method of claim 7, further comprising:
 periodically replicating state information among processes executed by the server computers to conduct the on-line conference;

15

 detecting a failure of one of the process;
 spawning a new process on the server computers; and

 loading the replicated state information on the new process.

20

9. The method of claim 8, wherein the processes whose state is replicated maintain information about the on-line conference.

25

10. The method of claim 8, wherein the processes whose state is replicated handle communications between one of the client computers and one of the server computers.

30

11. The method of claim 8, wherein the processes whose state is replicated control access to a document shared among participants of the on-line conference.

12. The method of claim 8, wherein the processes
whose state is replicated control execution of an
application shared among participants of the on-line
5 conference.

13. A computer-readable storage medium storing a
computer program executable by a plurality of server
computers, the computer program comprising computer
10 instructions for:

conducting an on-line conference among an
arbitrary number of the client computers connected
to an arbitrary number of the server computers via
the global-area network and the high-speed direct
15 connection link;

detecting a failure of one of the server
computers handling the on-line conference;

disconnecting the failed server computer from
the on-line conference;

20 connecting another of the server computers to
the conference; and

resuming the on-line conference.

14. The computer-readable storage medium of claim
25 13, wherein the computer program further comprises
computer instructions for:

periodically replicating state information
among processes executed by the server computers
to conduct the on-line conference;

30 detecting a failure of one of the process;
spawning a new process on the server
computers; and

[illegible]

18. The computer-readable storage medium of claim 14, wherein the processes whose state is replicated control execution of an application shared among participants of the on-line conference.